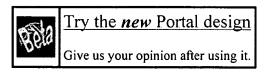


> home : > about : > feedback : > login

US Patent & Trademark Office



## Search Results

Search Results for: [(wraps around) AND ("tail" or "write pointer")] Found 61 of 129,763 searched.

Search within Results

GO > /

> Advanced Search

> Search Help/Tips

Sort by: Title Publication Publication Date Score Binder

Results 1 - 20 of 61

short listing

Prev Page 1 2 3

Circuits for wide-window superscalar processors

87%

Dana S. Henry , Bradley C. Kuszmaul , Gabriel H. Loh , Rahul Sami

ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture May 2000

Volume 28 Issue 2

Our program benchmarks and simulations of novel circuits indicate that large-window processors are feasible. Using our redesigned superscalar components, a large-window processor implemented in today's technology can achieve an increase of 10-60% (geometric mean of 31%) in program speed compared to today's processors. The processor operates at clock speeds comparable to today's processors, but achieves significantly higher ILP. To measure the impact of a large window on clock spe ...

Energy efficient architectures: Reducing power requirements of instruction scheduling through dynamic allocation of multiple datapath resources Dmitry Ponomarev , Gurhan Kucuk , Kanad Ghose

80%

**Proceedings of the 34th annual ACM/IEEE international symposium on Microarchitecture**December 2001

The "one-size-fits-all" philosophy used for permanently allocating datapath resources in today's superscalar CPUs to maximize performance across a wide range of applications results in the overcommitment of resources in general. To reduce power dissipation in the datapath, the resource allocations can be dynamically adjusted based on the demands of applications. We propose a mechanism to dynamically, simultaneously and independently adjust the sizes of the issue queue (IQ), the reorder buffer (R ...

Efficient use of memory bandwidth to improve network processor throughput Jahangir Hasan , Satish Chandra , T. N. Vijaykumar

80%

ACM SIGARCH Computer Architecture News, Proceedings of the 30th annual international symposium on Computer architecture May 2003
Volume 31 Issue 2

We consider the efficiency of packet buffers used in packet switches built using network processors (NPs). Packet buffers are typically implemented using DRAM, which provides plentiful buffering at a reasonable cost. The problem we address is that a typical NP workload may be unable to utilize the peak DRAM bandwidth. Since the bandwidth of the packet buffer is often the bottleneck in the performance of a shared-memory packet switch, inefficient use of available DRAM bandwidth further reduces th ...

Session 7: Tradeoffs in power-efficient issue queue design

80%

Alper Buyuktosunoglu , David H. Albonesi , Pradip Bose , Peter W. Cook , Stanley E. Schuster Proceedings of the 2002 international symposium on Low power electronics and design August 2002

A major consumer of microprocessor power is the issue queue. Several microprocessors, including the Alpha 21264 and POWER4™, use a compacting latch-based issue queue design which has the advantage of simplicity of design and verification. The disadvantage of this structure, however, is its high power dissipation. In this paper, we explore different issue queue power optimization techniques that vary not only in their performance and power characteristics, but in how much they deviate ...

Minimal adaptive routing with limited injection on Toroidal k-ary n-cubes

80%

Fabrizio Petrini, Marco Vanneschi

Proceedings of the 1996 ACM/IEEE conference on Supercomputing (CDROM) November 1996 Virtual channels can be used to implement deadlock free adaptive routing algorithms and increase network throughput. Unfortunately, they introduce asymmetries in the use of buffers of symmetric networks as the toroidal k-ary n-cubes. In this paper we present a minimal adaptive routing algorithm that tries to balance the use of the virtual channels by limiting the injection of new packets into the network. The experimental results, conducted on a 256 nodes torus, show that it is possible to ...

**6** An Optimal Algorithm for Finding the Kernel of a Polygon

80%

D. T. Lee , F. P. Preparata

Journal of the ACM (JACM) July 1979

Volume 26 Issue 3

Self-stabilization by window washing

80%

Adam M. Costello, George Varghese

Proceedings of the fifteenth annual ACM symposium on Principles of distributed computing May 1996

Extended ephemeral logging: log storage management for applications with long

80%

বী lived transactions John S. Keen , William J. Dally

ACM Transactions on Database Systems (TODS) March 1997

Volume 22 Issue 1

Numerical analysis using nonprocedural paradigms

80%

Stephen J. Sullivan, Benjamin G. Zorn

ACM Transactions on Mathematical Software (TOMS) September 1995

Volume 21 Issue 3

This article presents a survey on the innovative features of a handful of languages that offer new features that can be valuable in numerical analysis, and a survey of the pros and cons of the languages with regards to work in numerical analysis. Language features such as polymorphism, first-class functions, and object-oriented programming offer improved writability, readability,

reliability, and maintenance of computer software. The article discusses language features and uses, and include ...

10 An analysis of diffusive load-balancing

80%

Raghu Subramanian , Isaac D. Scherson

Proceedings of the sixth annual ACM symposium on Parallel algorithms and architectures August 1994

Diffusion is a well-known algorithm for load-balancing in which tasks move from heavily-loaded processors to lightly-loaded neighbors. This paper presents a rigorous analysis of the performance of the diffusion algorithm on arbitrary networks. It is shown that the running time of the diffusion algorithm is bounded by: &OHgr;(log &sgr;/&Ggr;) ≤ Time ≤ O(N&sgr;/&Ggr;) and &OHgr;(log \$sgr;/\$Fgr;)  $\le$  Time  $\le$  O(\$sgr;/\$Fgr;2), where N is the number of no ...

11 Performance evaluation of ephemeral logging

80%

John S. Keen , William J. Dally

ACM SIGMOD Record, Proceedings of the 1993 ACM SIGMOD international conference on Management of data June 1993

Volume 22 Issue 2

Ephemeral logging (EL) is a new technique for managing a log of database activity on disk. It does not require periodic checkpoints and does not abort lengthy transactions as frequently as traditional firewall logging for the same amount of disk space. Therefore, it is well suited for highly concurrent databases and applications which have a wide distribution of transaction lifetimes. This paper briefly explains EL and then analyzes its performance. Simulation studies indicate th ...

12 An architecture for optimal all-to-all personalized communication

80%

Susan Hinrichs, Corey Kosak, David R. O'Hallaron, Thomas M. Stricker, Riichiro Take Proceedings of the sixth annual ACM symposium on Parallel algorithms and architectures August 1994

In all-to-all personalized communication (AAPC), every node of a parallel system sends a potentially unique packet to every other node. AAPC is an important primitive operation for modern parallel compilers, since it is used to redistribute data structures during parallel computations. As an extremely dense communication pattern, AAPC causes congestion in many types of networks and therefore executes very poorly on general purpose, asynchronous message passsing routers. We presen ...

13 An efficient architecture for loop based data preloading

80%

William Y. Chen , Roger A. Bringmann , Scott A. Mahlke , Richard E. Hank , James E. Sicolo ACM SIGMICRO Newsletter, Proceedings of the 25th annual international symposium on Microarchitecture December 1992

Volume 23 Issue 1-2

14 Design and verification of the Rollback Chip using HOP: a case study of formal

80%

nethods applied to hardware design Ganesh Gopalakrishnan, Richard Fujimoto

ACM Transactions on Computer Systems (TOCS) May 1993

Volume 11 Issue 2

The use of formal methods in hardware design improves the quality of designs in many ways: it promotes better understanding of the design; it permits systematic design refinement through the discovery of invariants; and it allows design verification (informal or formal). In this paper we illustrate the use of formal methods in the design of a custom hardware system called the "Rollback Chip" (RBC), conducted using a simple hardware design description language called "HOP&r ...

**15** The performance analysis workstation: an interactive animated simulation package

বী for queueing networks

B. Melamed

Proceedings of 1986 fall joint computer conference on Fall joint computer conference November 1999

**16** Emerging areas: A new look at exploiting data parallelism in embedded systems

77%

Hillery C. Hunter , Jaime H. Moreno

Proceedings of the international conference on Compilers, architectures and synthesis for embedded systems October 2003

This paper describes and evaluates three architectural methods for accomplishing data parallel computation in a programmable embedded system. Comparisons are made between the well-studied Very Long Instruction Word (VLIW) and Single Instruction Multiple Packed Data (SIMpD) paradigms; the less-common Single Instruction Multiple Disjoint Data (SIMdD) architecture is described and evaluated. A taxonomy is defined for data-level parallel archi ...

**17** VMTP: a transport protocol for the next generation of communication systems

77%

D Cheriton

Proceedings of the ACM SIGCOMM conference on Communications architectures & protocols September 1986

The Versatile Message Transaction Protocol (VMTP) is a transport-level protocol designed to support remote procedure call, multicast and real-time communication. The protocol is optimized for efficientpage-level network file access in particular. In this paper, we describe the significant aspects of the VMTP design, including the VMTP treatment of sessions, addressing, duplicate suppression, flow control and retransmissions plus its provision for multicast. The VMTP design refle ...

**18** PERUSE: An Interactive System for Mathematical Programs

77%

William G. Kurator , Richard P. O'Neill

ACM Transactions on Mathematical Software (TOMS) December 1980 Volume 6 Issue 4

**19** Applications II: Time and area efficient pattern matching on FPGAs

77%

Zachary K. Baker , Viktor K. Prasanna

Proceeding of the 2004 ACM/SIGDA 12th international symposium on Field programmable gate arrays February 2004

Pattern matching for network security and intrusion detection demands exceptionally high performance. Much work has been done in this field, and yet there is still significant room for improvement in efficiency, flexibility, and throughput. We develop a novel linear-array string matching architecture using a buffered, two-comparator variation on the Knuth-Morris-Pratt(KMP) algorithm. For small (16 or fewer characters) patterns, it competes favorably with the state-of-theart while providing bett ...

20 Transputers + virtual tree kernel = real speedups

77%

D. McBurney, M. R. Sleep

Proceedings of the third conference on Hypercube concurrent computers and applications: Architecture, software, computer systems, and general issues - Volume 1 January 1988 In [9] we describe a simple virtual tree architecture intended to exploit parallelism in recursive algorithms. The architecture, called ZAPP (Zero Assignment Parallel Processor), is relatively simple to implement using transputers. Early experiments with small configurations were encouraging both with respect to relative speedups and, somewhat to our surprise as ZAPP was simulated, absolute performance. In this paper we report more recent experiments with matrix multiplication an ...

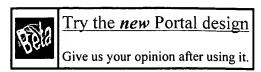
Results 1 - 20 of 61

short listing





> home : > about : > feedback US Patent & Trademark Office



## Search Results

Search Results for: [(wraps around) AND ("tail" or "write pointer")] Found **61** of **129,763 searched.** 

Search within Results

GO

> Advanced Search : > Search Help/Tips

Binder Sort by: **Title Publication Publication Date** Score

Results 21 - 40 of 61 short listing

21 Virtual simple architecture (VISA): exceeding the complexity limit in safe real-time 77% ৰী systems

Aravindh Anantaraman, Kiran Seth, Kaustubh Patil, Eric Rotenberg, Frank Mueller ACM SIGARCH Computer Architecture News, Proceedings of the 30th annual international symposium on Computer architecture May 2003

Volume 31 Issue 2

Meeting deadlines is a key requirement in safe realtime systems. Worst-case execution times (WCET) of tasks are needed for safe planning. Contemporary worst-case timing analysis tools can safely and tightly bound execution time on in-order single-issue pipelines with caches and static branch prediction. However, this simple pipeline appears to be a complexity limit, due to the need for analyzability. This excludes a whole class of high-performance processors from many embedded systems. We reconci ...

22 Synthesis of saturation arithmetic architectures

77%

G. A. Constantinides , P. Y. K. Cheung , W. Luk

ACM Transactions on Design Automation of Electronic Systems (TODAES) July 2003 Volume 8 Issue 3

This paper describes a synthesis technique for automating the design of linear Digital Signal Processing (DSP) systems such as digital filters. The proposed methodology makes optimized use of saturation arithmetic to produce a small design implemented directly in hardware. An analytical technique is proposed to estimate the saturation error resulting from a particular implementation, and an optimization procedure is introduced to aim for the smallest implementation satisfying userspecified boun ...

23 Distance-based outliers: algorithms and applications

77%

Edwin M. Knorr, Raymond T. Ng, Vladimir Tucakov

The VLDB Journal — The International Journal on Very Large Data Bases February 2000 Volume 8 Issue 3-4

This paper deals with finding outliers (exceptions) in large, multidimensional datasets. The

identification of outliers can lead to the discovery of truly unexpected knowledge in areas such as electronic commerce, credit card fraud, and even the analysis of performance statistics of professional athletes. Existing methods that we have seen for finding outliers can only deal efficiently with two dimensions/attributes of a dataset. In this paper, we study the notion of DB (distance-based ...

**24** K9: a simulator of distributed-memory parallel processors

77%

P. Beadle , C. Pommerell , M. Annaratone

Proceedings of the 1989 ACM/IEEE conference on Supercomputing August 1989

K9 is a software package for the simulation and performance evaluation of distributed-memory parallel processors (DMPPs). It is written in C++ and runs on Sequent Symmetry and SUN-3. K9 provides the user with four building-blocks (processor cells, communication channels, multi-port shared-memories, and I/O processors), and one abstraction mechanism (the DMPP interconnection topology). Application code for K9 can be written in C++ or C. When timing analysi ...

25 Multiprocessors with a serial multiport memory and a pseudo crossbar of serial links 77%

বী used s a processor-memeory switch

Daniel Litaize, Omar Hammami, Mustapha Lalam, Adelaziz Mzoughi, Pascl Sinrat **ACM SIGARCH Computer Architecture News** December 1989

Volume 17 Issue 6

This paper presents an inventive information exchange pro-cess between the main memory and cache equipped processors. It makes use of serial multiport memories and high throughput serial transmission supports. It is then possible to consider the realization of a multiprocessor with a common memory shared by several hundreds processors set with a performance level close to that of a crossbar network one's without having its disadvantages. This exchange process generates a family of possible archi ...

**26** High performance communications in processor networks

77%

🔥 C. R. Jesshope , P. R. Miller , J. T. Yantchev

ACM SIGARCH Computer Architecture News, Proceedings of the 16th annual international symposium on Computer architecture April 1989

Volume 17 Issue 3

In order to provide an arbitrary and fully dynamic connectivity in a network of processors, transport mechanisms must be implemented, which provide the propagation of data from processor to processor, based on addresses contained within a packet of data. Such data transport mechanisms must satisfy a number of requirements - deadlock and livelock freedom, good hot-spot performance, high throughput and low latency. This paper proposes a solution to these problems, which allows deadlock free, ...

27 A stream compiler for communication-exposed architectures

77%

Michael I. Gordon, William Thies, Michal Karczmarek, Jasper Lin, Ali S. Meli, Andrew A. Lamb, Chris Leger, Jeremy Wong, Henry Hoffmann, David Maze, Saman Amarasinghe

Tenth international conference on architectural support for programming languages and operating systems on Proceedings of the 10th international conference on architectural support for programming languages and operating systems (ASPLOS-X) October 2002 Volume 37, 30, 36 Issue 10, 5, 5

With the increasing miniaturization of transistors, wire delays are becoming a dominant factor in microprocessor performance. To address this issue, a number of emerging architectures contain replicated processing units with software-exposed communication between one unit and another (e.g., Raw, SmartMemories, TRIPS). However, for their use to be widespread, it will be necessary to develop compiler technology that enables a portable, high-level language to execute efficiently across a range of w ...

**28** Data buffer performance for sequential Prolog architectures

77%

ACM SIGARCH Computer Architecture News, Proceedings of the 15th Annual International Symposium on Computer architecture May 1988

Volume 16 Issue 2

Several local data buffers are proposed and measurements are presented for variations of the Warren Abstract Machine (WAM) architecture for Prolog. Choice point buffers, stack buffers, splitstack buffers, multiple register sets, copyback caches, and "smart" caches are examined. Statistics collected from four benchmark programs indicate that small conventional local memories perform quite well because of the WAM's high locality. The data memory performance results are equally va ...

29 Transaction papers: A dynamic call admission policy with precision QoS quarantee

77%

बो using stochastic control for mobile wireless networks Si Wu, K. Y. Michael Wong, Bo Li

IEEE/ACM Transactions on Networking (TON) April 2002

Volume 10 Issue 2

Call admission control is one of the key elements in ensuring the quality of serivce in mobile wireless networks. The traditional trunk reservation policy and its numerous variants give preferential treatment to the handoff calls over new arrivals by reserving a number of radio channels exclusively for handoffs. Such schemes, however, cannot adapt to changes in traffic pattern due to the static nature. This paper introduces a novel stable dynamic call admission control mechanism (SDCA), which ca ...

**30** Transaction papers: Formal specification and verification of safety and performance .77%

4 of TCP selective acknowledgment Mark A. Smith , K. K. Ramakrishnan

IEEE/ACM Transactions on Networking (TON) April 2002

Volume 10 Issue 2

We present a formal specification of the selective acknowledgment (SACK) mechanism that is being proposed as a new standard option for TCP. The formal specification allows one to reason about the SACK protocol; thus, we are able to formally prove that the SACK mechanism does not violate the safety properties (reliable, at most once, and in order message delivery) of the acknowledgment (ACK) mechanism that is currently used with TCP. The new mechanism is being proposed to improve the performance ...

**31** Optimizing the end-to-end performance of reliable flows over wireless links Reiner Ludwig , Almudena Konrad , Anthony D. Joseph , Randy H. Katz

77%

Wireless Networks March 2002

Volume 8 Issue 2/3

Pure end-to-end error recovery fails as a general solution to optimize throughput when wireless links form parts of the end-to-end path. It can lead to decreased end-to-end throughput, an unfair load on best-effort networks, and a waste of valuable radio resources. Link layer error recovery over wireless links is essential for reliable flows to avoid these problems. We demonstrate this through an analysis of a large set of block erasure traces measured in different real-world radio environments,

**32** Guard against data loss with mondo rescue

77%

Hugo Rabson

Linux Journal December 2001

Volume 2001 Issue 92

Looking for an easy open-source backup method?

**33** Credit-based fair queueing (CBFQ): a simple service-scheduling algorithm for

77%

| | packet-switched networks

Brahim Bensaou, Danny H. K. Tsang, King Tung Chan

IEEE/ACM Transactions on Networking (TON) October 2001

Volume 9 Issue 5

This paper proposes a simple rate-based scheduling algorithm for packet-switched networks. Using a set of counters to keep track of the credits accumulated by each traffic flow, the bandwidth share allocated to each flow, and the size of the head-of-line (HOL) packets of the different flows, the algorithm decides which flow to serve next. Our proposed algorithm requires on average a smaller complexity than the most interesting alternative ones while guaranteeing comparable fairness, delay, and d ...

**34** Illustrating programmed and interrupt driven I/O

77%

Terry A. Scott

The Journal of Computing in Small Colleges October 2000

Volume 16 Issue 1

**35** Diffusion tree restructuring for indirect reference counting

77%

Peter Dickman

ACM SIGPLAN Notices, Proceedings of the second international symposium on Memory management October 2000

Volume 36 Issue 1

A new variant algorithm for distributed acyclic garbage detection is presented for use in hybrid garbage collectors. The existing fault-tolerance of Piquer's Indirect Reference Counting (IRC) is qualitatively improved by this new approach. The key insight that underpins this work is the observation that the parent of a node in the IRC diffusion tree need not remain constant. The new variant exploits standard mechanisms for implementing diffusion trees and remote references, using four simple ...

**36** Cheap eagerness: speculative evaluation in a lazy functional language

77%

Karl-Filip Faxén

ACM SIGPLAN Notices, Proceedings of the fifth ACM SIGPLAN international conference on Functional programming September 2000

Volume 35 Issue 9

Cheap eagerness is an optimization where cheap and safe expressions are evaluated before it is known that their values are needed. Many compilers for lazy functional languages implement this optimization, but they are limited by a lack of information about the global flow of control and about which variables are already evaluated. Without this information, even a variable reference is a potentially unsafe expression!In this paper we show that significant speedups are achievable by cheap eagernes ...

**37** Recovery management in QuickSilver

77%

Rober Haskin , Yoni Malachi , Gregory Chan

ACM Transactions on Computer Systems (TOCS) February 1988

Volume 6 Issue 1

This paper describes QuickSilver, developed at the IBM Almaden Research Center, which uses atomic transactions as a unified failure recovery mechanism for a client-server structured distributed system. Transactions allow failure atomicity for related activities at a single server or at a number of independent servers. Rather than bundling transaction management into a dedicated language or recoverable object manager, Quicksilver exposes the basic commit protocol and log rec ...

**38** Improving parallel system performance by changing the arrangement of the

77%

network links

V. Puente, C. Izu, J. A. Gregorio, R. Beivide, J. M. Prellezo, F. Vallejo

Proceedings of the 14th international conference on Supercomputing May 2000

The Midimew network is an excellent contender for implementing the communication subsystem of a high performance computer. This network is an optimal 2D topology in the sense there are no other symmetric direct networks of degree 4 with a lower average distance or diameter. In fact, it reduces the diameter of the well known torus network by approximately □ 2. Although the topology was proposed and analyzed a decade ago, the lack of simple deadlock avoidance mechanisms prevented its ut ...

**39** System-level power optimization: techniques and tools

77%

🖈 Luca Benini , Giovanni de Micheli

ACM Transactions on Design Automation of Electronic Systems (TODAES) April 2000 Volume 5 Issue 2

This tutorial surveys design methods for energy-efficient system-level design. We consider electronic sytems consisting of a hardware platform and software layers. We consider the three major constituents of hardware that consume energy, namely computation, communication, and storage units, and we review methods of reducing their energy consumption. We also study models for analyzing the energy cost of software, and methods for energy-efficient software design and compilation. This survery ...

**40** An empirical study on how program layout affects cache miss rates

77%

Jeffrey P. Bradford , Russell Quona

**ACM SIGMETRICS Performance Evaluation Review** December 1999

Volume 27 Issue 3

Cache miss rates are quoted for a specific program, cache configuration, and input set; the effect of program layout on the miss rate has largely been ignored. This paper examines the miss variation, that is, the variation in the miss rate for instruction and data caches resulting from randomly generated layouts; the layouts were generated by changing the order of the modules on the command line when linking. This analysis is performed for several cache sizes, lines sizes, setassociat ...

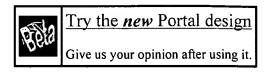
Results 21 - 40 of 61

short listing





> home : > about : > feedback **US Patent & Trademark Office** 



Search Results

Search Results for: [(wraps around) AND ("tail" or "write pointer")] Found **61** of **129,763 searched.** 

Search within Results

> Advanced Search

> Search Help/Tips

Sort by: **Title** 

**Publication Publication Date**  Score

GO

Binder

Results 41 - 60 of 61

short listing

**41** The effects of asymmetry on TCP performance

Hari Balakrishnan , Randy H. Katz , Venkata N. Padmanbhan Mobile Networks and Applications October 1999

Volume 4 Issue 3

In this paper, we study the effects of network asymmetry on end-to-end TCP performance and suggest techniques to improve it. The networks investigated in this study include a wireless cable modem network and a packet radio network, both of which can form an important part of a mobile ad hoc network. In recent literature (e.g., [18]), asymmetry has been considered in terms of a mismatch in bandwidths in the two directions of a data transfer. We generalize this notion of bandwidth asymmetry t ...

**42** A type system for expressive security policies

77%

David Walker

Proceedings of the 27th ACM SIGPLAN-SIGACT symposium on Principles of programming languages January 2000

Certified code is a general mechanism for enforcing security properties. In this paradigm, untrusted mobile code carries annotations that allow a host to verify its trustworthiness. Before running the agent, the host checks the annotations and proves that they imply the host's security policy. Despite the flexibility of this scheme, so far, compilers that generate certified code have focused on simple type safety properties rather than more general security properties.

**43** Ace: a language for parallel programming with customizable protocols

77%

Mukund Raghavachari, Anne Rogers

ACM Transactions on Computer Systems (TOCS) August 1999

Volume 17 Issue 3

Customizing the protocols that manage accesses to different data structures within an application can improve the performance of software shared-memory programs substantially. Existing systems for using customizable protocols are hard to use directly because the mechanisms they provide for manipulating protocols are low-level ones. This article is an in-depth study of the issues involved in providing language support for application-specific protocols. We describe the design and

implementat ...

44 Coverage estimation for symbolic model checking

77%

Yatin Hoskote , Timothy Kam , Pei-Hsin Ho , Xudong Zhao

Proceedings of the 36th ACM/IEEE conference on Design automation conference June 1999

**45** A comparison of the concurrency features of Ada 95 and Java

77%

Benjamin M. Brosgol

ACM SIGAda Ada Letters, Proceedings of the 1998 annual ACM SIGAda international conference on Ada November 1998

Volume XVIII Issue 6

**46** Timestamp representations for virtual sequences

77%

John G. Cleary , J. A. David McWha , Murray Pearson

ACM SIGSIM Simulation Digest, Proceedings of the eleventh workshop on Parallel and distributed simulation June 1997

Volume 27 Issue 1

The problem of executing sequential programs in parallel using the optimistic algorithm Time Warp is considered. This is done by first mapping the sequential execution to a control tree and then assigning timestamps to each node in the tree. For such timestamps to be effective in either hardware or software they must be finite, this implies that they must be periodically rescaled to allow old timestamps to be reused. A number of timestamp representations are described and compared on the basis of ...

47 Teaching ethical and social issues in CS1 and CS2

77%

Kay G. Schulze , Frances S. Grodzinsky

ACM SIGCSE Bulletin, Proceedings of the twenty-eighth SIGCSE technical symposium on Computer science education March 1997

Volume 29 Issue 1

The discussion of whether ethical and social issues of computing should be explored in undergraduate computer science education has resulted in most academic institutions and educators agreeing that they are important topics that must be included. Further support has been provided by Curricula '91 [16], the CSAC/CSAB accreditation [2] and ImpactCS [12]. Many books [7, 8, 9, 10] and papers [6, 14] have discussed what topics should be covered and what techniques can be used either in a dedicated c ...

**48** On tuning the microarchitecture of an HPS implementation of the VAX

77%

James E. Wilson , Steve Melvin , Michael Shebanow , Wen-mei Hwu , Yale N. Patt

Proceedings of the 20th annual workshop on Microprogramming December 1987 The HPS Microarchitecture has been developed as an execution model for implementing various architectures at very high performance. A considerable amount of effort has gone into the use of HPS as a microarchitecture for the VAX. In this paper, we describe our first full simulation of the microVAX subset, and report the results of varying (i.e. tuning) certain important parameters.

49 On the benefit of supporting virtual channels in wormhole routers

77%

Richard J. Cole , Bruce M. Maggs , Ramesh K. Sitaraman

Proceedings of the eighth annual ACM symposium on Parallel algorithms and architectures June 1996

**50** Cube-4—a scalable architecture for real-time volume rendering Hanspeter Pfister, Arie Kaufman

77%

Proceedings of the 1996 symposium on Volume visualization October 1996

**51** Coherent network interfaces for fine-grain communication

77%

Shubhendu S. Mukherjee , Babak Falsafi , Mark D. Hill , David A. Wood

ACM SIGARCH Computer Architecture News, Proceedings of the 23rd annual international symposium on Computer architecture May 1996

Volume 24 Issue 2

Historically, processor accesses to memory-mapped device registers have been marked uncachable to insure their visibility to the device. The ubiquity of snooping cache coherence, however, makes it possible for processors and devices to interact with cachable, coherent memory operations. Using coherence can improve performance by facilitating burst transfers of whole cache blocks and reducing control overheads (e.g., for polling). This paper begins an exploration of network interfaces (NIs) that u ...

**52** Dynamic self-invalidation: reducing coherence overhead in shared-memory

77%

multiprocessors

Alvin R. Lebeck , David A. Wood

ACM SIGARCH Computer Architecture News, Proceedings of the 22nd annual international symposium on Computer architecture May 1995

Volume 23 Issue 2

This paper introduces dynamic self-invalidation (DSI), a new technique for reducing cache coherence overhead in shared-memory multiprocessors. DSI eliminates invalidation messages by having a processor automatically invalidate its local copy of a cache block before a conflicting access by another processor. Eliminating invalidation overhead is particularly important under sequential consistency, where the latency of invalidating outstanding copies can increase a program's critical path.DSI is ap ...

**53** A measurement-based admission control algorithm for integrated services packet

77%

বি networks

Sugih Jamin, Peter B. Danzig, Scott Shenker, Lixia Zhang

ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication October 1995

Volume 25 Issue 4

Many designs for integrated service networks offer a bounded delay packet delivery service to support real-time applications. To provide bounded delay service, networks must use admission control to regulate their load. Previous work on admission control mainly focused on algorithms that compute the worst case theoretical queueing delay to guarantee an absolute delay bound for all packets. In this paper we describe a measurement-based admission control algorithm for predictive serv ...

**54** Diversity-based inference of finite automata

77%

Ronald L. Rivest, Robert E. Schapire

Journal of the ACM (JACM) May 1994

Volume 41 Issue 3

We present new procedures for inferring the structure of a finite-state automaton (FSA) from its input/output behavior, using access to the automaton to perform experiments. Our procedures use a new representation for finite automata, based on the notion of equivalence between tests. We call the number of such equivalence classes the diversity of the automaton; the diversity may be as small as the logarithm of the number of states of the automato ...

**55** Enhanced superscalar hardware: the schedule table

77%

J. K. Pickett , D. G. Meyer

Proceedings of the 1993 ACM/IEEE conference on Supercomputing December 1993

**56** Lightweight recoverable virtual memory

77%

M. Satyanarayanan , Henry H. Mashburn , Puneet Kumar , David C. Steere , James J. Kistler ACM SIGOPS Operating Systems Review, Proceedings of the fourteenth ACM symposium on Operating systems principles December 1993

Volume 27 Issue 5

Recoverable virtual memory refers to regions of a virtual address space on which transactional guarantees are offered. This paper describes RVM, an efficient, portable, and easily used implementation of recoverable virtual memory for Unix environments. A unique characteristic of RVM is that it allows independent control over the transactional properties of atomicity, permanence, and serializability. This leads to considerable flexibility in the use of RVM, potentially enlarging the ...

**57** Analysis, modeling and generation of self-similar VBR video traffic

77%

Mark W. Garrett , Walter Willinger

ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Communications architectures, protocols and applications October 1994

Volume 24 Issue 4

We present a detailed statistical analysis of a 2-hour long empirical sample of VBR video. The sample was obtained by applying a simple intraframe video compression code to an action movie. The main findings of our analysis are (1) the tail behavior of the marginal bandwidth distribution can be accurately described using "heavy-tailed" distributions (e.g., Pareto); (2) the autocorrelation of the VBR video sequence decays hyperbolically (equivalent to long-range dependenc ...

**58** An empirical study of a highly available file system

77%

Brian D. Noble, M. Satyanarayanan

ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1994 ACM SIGMETRICS conference on Measurement and modeling of computer systems May 1994 Volume 22 Issue 1

In this paper we present results from a six-month empirical study of the high availability aspects of the Coda File System. We report on the service failures experienced by Coda clients, and show that such failures are masked successfully. We also explore the effectiveness and resource costs of key aspects of server replication and disconnected operation, the two high availability mechanisms of Coda. Wherever possible, we compare our measurements to simulat ...

**59** Improving instruction supply efficiency in superscalar architectures using instruction 77%

বী trace buffers Chih-Po Wen

> Proceedings of the 1992 ACM/SIGAPP Symposium on Applied computing: technological challenges of the 1990's April 1992

**60** Planar graph decomposition and all pairs shortest paths

77%

Grea N. Frederickson

Journal of the ACM (JACM) January 1991

Volume 38 Issue 1

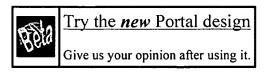
An algorithm is presented for generating a succinct encoding of all pairs shortest path information in a directed planar graph G with real-valued edge costs but no negative cycles. The algorithm runs in O(pn) time, where n is the number of vertices in G, and p is the minimum cardinality of a subset of the faces that cover all vertices, taken over all planar embeddings of G. ...

Results 41 - 60 of 61 short listing





> home : > about : > feedback US Patent & Trademark Office



Search Results

Search Results for: [(wraps around) AND ("tail" or "write pointer")] Found **61** of **129,763 searched.** 

Search within Results

GO

> Advanced Search

> Search Help/Tips

Binder Sort by: Title **Publication Publication Date** Score

Results 61 - 61 of 61 short listing

**61** A stepwise refinement heuristic for protocol construction

77%

A. Udaya Shankar, Simon S. Lam

ACM Transactions on Programming Languages and Systems (TOPLAS) May 1992

Volume 14 Issue 3

A stepwise refinement heuristic to construct distributed systems is presented. The heuristic is based on a conditional refinement relation between system specifications, and a "Marking". It is applied to construct four sliding window protocols that provide reliable data transfer over unreliable communication channels. The protocols use modulo-N sequence numbers. The first protocol is for channels that can only lose messages in transit. By refining this protocol, ...

Results 61 - 61 of 61

short listing

